In the claims

1. (currently amended) A method comprising:

mapping a plurality of physically non-contiguous sections of memory into a logically contiguous section of memory; ~~and,~~

for each computing unit of a plurality of computing units, allocating a portion of the logically contiguous section of memory addressable by a pointer plus a static offset corresponding to the computing unit,

wherein the static offset for each computing unit is equal to a static offset initially determined at initial allocation of memory for the plurality of computing units;

dynamically passing out the portion of the logically contiguous section of memory to each computing unit of the plurality of computing units as the computing units need additional memory;

upon the logically contiguous section of memory being completely passed out to the plurality of computing units,

mapping a second plurality of physically non-contiguous sections of memory into a second logically contiguous section of memory;

for each computing unit of the plurality of computing units, allocating a portion of the second logically contiguous section of memory addressable by a pointer plus the static offset corresponding to the computing unit; and,

dynamically passing out the portion of the second logically contiguous section of memory to each computing unit of the plurality of computing units as the computing units need additional memory.

2. (original) The method of claim 1, wherein the portion of the logically contiguous section of memory allocated for each computing unit includes memory local to the computing unit.

3. (cancelled)

4. (original) The method of claim 1, further comprising determining the static offset for each computing unit as equal to the static offset initially determined at the initial allocation of the memory for the plurality of computing units.

5. (original) The method of claim 1, further comprising at the initial allocation of the memory for the plurality of computing units:

determining the static offset for each computing unit of the plurality of computing units; and,

for each computing unit of the plurality of computing units, allocating a portion of memory addressable by a pointer plus the static offset corresponding to the computing unit.

6. (original) The method of claim 5, further comprising dynamically passing out the portion of the memory to each computing unit of the plurality of computing units as the computing units need additional memory.

7. (original) The method of claim 1, wherein the computing unit is one of a computing node and a processor.

8. (original) A method comprising:

determining a static offset for each computing unit of a plurality of computing units;

for each computing unit of the plurality of computing units, allocating a portion of predetermined memory addressable by a pointer plus the static offset corresponding to the computing unit;

dynamically passing out the portion of the predetermined memory to each computing unit of the plurality of computing units as the computing units need additional memory;

upon the predetermined memory being completely passed out to the plurality of computing units,

mapping a plurality of physically non-contiguous sections of additional memory into a logically contiguous section of memory;

for each computing unit of the plurality of computing units, allocating a portion of the logically contiguous section of memory addressable by a pointer plus the static offset corresponding to the computing unit, where the portion of the logically contiguous section of memory is local to the computing unit; and,

dynamically passing out the portion of the logically contiguous section of memory to each computing unit of the plurality of computing units as the computing units need additional memory.

9. (original) The method of claim 8, further comprising, upon the logically contiguous section of memory being completely passed out to the plurality of computing units,

mapping a second plurality of physically non-contiguous sections of memory into a second logically contiguous section of memory;

for each computing unit of the plurality of computing units, allocation a portion of the second logically contiguous section of memory addressable by a pointer plus the static offset corresponding to the computing unit, wherein the portion of the second logically contiguous section of memory is local to the computing unit; and,

dynamically passing out the portion of the second logically contiguous section of memory to each computing unit of the plurality of computing units as the computing units needs additional memory.

10. (currently amended) A system comprising:

a plurality of computing nodes;

memory shared by the plurality of computing nodes; and,

an allocating mechanism to map a plurality of physically non-contiguous sections of memory into a logically contiguous section of memory and to allocate a portion of the logically contiguous section of memory to each computing node,

wherein the portion of the logically contiguous section of memory is addressable by a pointer plus a static offset corresponding to the computing node and equal to a static offset initially determined for allocating a portion of memory to each computing node, and

wherein the allocating mechanism is further to

dynamically pass out the portion of the logically contiguous section of memory to each computing node of the plurality of computing units as the computing nodes need additional memory;

upon the logically contiguous section of memory being completely passed out to the plurality of computing nodes,

map a second plurality of physically non-contiguous sections of memory into a second logically contiguous section of memory;

for each computing node of the plurality of computing nodes, allocate a portion of the second logically contiguous section of memory addressable by a pointer plus the static offset corresponding to the computing node; and,

dynamically pass out the portion of the second logically contiguous section of memory to each computing node of the plurality of computing units as the computing nodes need additional memory.

11. (original) The system of claim 10, wherein the portion of the logically contiguous section of memory allocated to each computing node includes memory local to the computing node.

12. (original) The system of claim 10, wherein each of the plurality of computing nodes comprises a single processor.

13. (original) The system of claim 10, wherein each of the plurality of computing nodes comprises a plurality of processors.

14. (currently amended) A system comprising:

a plurality of computing nodes;

memory shared by the plurality of computing nodes; and,

means for mapping a plurality of physically non-contiguous sections of memory into a logically contiguous section of memory and for allocating a portion of the logically contiguous section of memory to each computing node,

wherein the portion of the logically contiguous section of memory is addressable by a pointer plus a static offset corresponding to the computing node and equal to a static offset for allocating a portion of memory to each computing node, and

wherein the means is further for

dynamically passing out the portion of the logically contiguous section of memory to each computing node of the plurality of computing units as the computing nodes need additional memory;

upon the logically contiguous section of memory being completely passed out to the plurality of computing nodes,

mapping a second plurality of physically non-contiguous sections of memory into a second logically contiguous section of memory;

for each computing node of the plurality of computing nodes, allocating a portion of the second logically contiguous section of memory addressable by a pointer plus the static offset corresponding to the computing node; and,

dynamically passing out the portion of the second logically contiguous section of memory to each computing node of the plurality of computing units as the computing nodes need additional memory.

15. (currently amended) A computing node comprising:

a plurality of processors;

memory shared by the plurality of processors; and,

an allocating mechanism to map a plurality of physically non-contiguous sections of memory into a logically contiguous section of memory and to allocate a portion of the logically contiguous section of memory to each processor,

wherein the portion of the logically contiguous section of memory is addressable by a pointer plus a static offset corresponding to the processor and equal to a static offset initially determined for allocating a portion of memory to each processor, and

wherein the allocating mechanism is further to

dynamically pass out the portion of the logically contiguous section of memory to each processor of the plurality of computing units as the processors need additional memory;

upon the logically contiguous section of memory being completely passed out to the plurality of processors,

map a second plurality of physically non-contiguous sections of memory into a second logically contiguous section of memory;

for each processor of the plurality of processors, allocate a portion of the second logically contiguous section of memory addressable by a pointer plus the static offset corresponding to the processor; and,

dynamically pass out the portion of the second logically contiguous section of memory to each processor of the plurality of computing units as the processors need additional memory.

16. (original) The computing node of claim 15, wherein the portion of the logically contiguous section of memory allocated to each processor is local to the processor.

17. (currently amended) A computing node comprising:

a plurality of processors;

memory shared by the plurality of processors; and,

means for mapping a plurality of physically non-contiguous sections of memory into a logically contiguous section of memory and for allocating a portion of the logically contiguous section of memory to each processor,

wherein the portion of the logically contiguous section of memory is addressable by a pointer plus a static offset corresponding to the processor and equal to a static offset for allocating a portion of memory to each processor, and

wherein the means is further for

dynamically passing out the portion of the logically contiguous section of memory to each processor of the plurality of computing units as the processors need additional memory;

upon the logically contiguous section of memory being completely passed out to the plurality of processors,

mapping a second plurality of physically non-contiguous sections of memory into a second logically contiguous section of memory;

for each processor of the plurality of processors, allocating a portion of the second logically contiguous section of memory addressable by a pointer plus the static offset corresponding to the processor; and,

dynamically passing out the portion of the second logically contiguous section of memory to each processor of the plurality of computing units as the processors need additional memory.

18. (currently amended) An article of manufacture comprising:

a computer-readable medium; and,

means in the medium for mapping a plurality of physically non-contiguous sections of memory into a logically contiguous section of memory and for allocating a portion of the logically contiguous section of memory to each computing unit of a plurality of computing units,

wherein the portion of the logically contiguous section of memory is addressable by a pointer plus a static offset corresponding to the computing unit and equal to a static offset initially determined at boot time of the plurality of computing units for allocating a portion of memory to each computing unit, and

wherein the means is further for

dynamically passing out the portion of the logically contiguous section of memory to each computing unit of the plurality of computing units as the computing units need additional memory;

upon the logically contiguous section of memory being completely passed out to the plurality of computing units,

mapping a second plurality of physically non-contiguous sections of memory into a second logically contiguous section of memory;

for each computing unit of the plurality of computing units, allocating a portion of the second logically contiguous section of memory addressable by a pointer plus the static offset corresponding to the computing unit; and,

dynamically passing out the portion of the second logically contiguous section of memory to each computing unit of the plurality of computing units as the computing units need additional memory.